

Unified Acceleration Method for Packing and Covering Problems via Diameter Reduction

Di Wang *

Satish Rao †

Michael W. Mahoney ‡

Abstract

The linear coupling method was introduced recently by Allen-Zhu and Orecchia [14] for solving convex optimization problems with first order methods, and it provides a conceptually simple way to integrate a gradient descent step and mirror descent step in each iteration. In the setting of standard smooth convex optimization, the method achieves the same convergence rate as that of the accelerated gradient descent method of Nesterov [8]. The high-level approach of the linear coupling method is very flexible, and it has shown initial promise by providing improved algorithms for packing and covering linear programs [1, 2]. Somewhat surprisingly, however, while the dependence of the convergence rate on the error parameter ϵ for packing problems was improved to $O(1/\epsilon)$, which corresponds to what accelerated gradient methods are designed to achieve, the dependence for covering problems was only improved to $O(1/\epsilon^{1.5})$, and even that required a different more complicated algorithm. Given the close connections between packing and covering problems and since previous algorithms for these very related problems have led to the same ϵ dependence, this discrepancy is surprising, and it leaves open the question of the exact role that the linear coupling is playing in coordinating the complementary gradient and mirror descent step of the algorithm. In this paper, we clarify these issues for linear coupling algorithms for packing and covering linear programs, illustrating that the linear coupling method can lead to improved $O(1/\epsilon)$ dependence for both packing and covering problems in a unified manner, i.e., with the same algorithm and almost identical analysis. Our main technical result is a novel diameter reduction method for covering problems that is of independent interest and that may be useful in applying the accelerated linear coupling method to other combinatorial problems.

1 Introduction

A fractional covering problem, in its generic form, can be written as the following linear program (LP):

$$\min_{x \geq 0} \{c^T x : Ax \geq b\},$$

where $c \in \mathbb{R}_{\geq 0}^n$, $b \in \mathbb{R}_{\geq 0}^m$, and $A \in \mathbb{R}_{\geq 0}^{m \times n}$. That is, we want to put weights on the x_i -s, for $i \in \{1, \dots, n\}$, such that each $j \in \{1, \dots, m\}$ is “covered” with weight at least b_j , where each unit of weight on x_i puts A_{ij} weight on each j , and we want to minimize the cost $c^T x$ in doing so. Without loss of generality, one can scale the coefficients, in which case one can write this LP in the standard form:

$$\min_{x \geq 0} \{\bar{1}^T x : Ax \geq \bar{1}\}, \tag{1}$$

where $A \in \mathbb{R}_{\geq 0}^{m \times n}$. The dual of this LP, the fractional packing problem, can be written in this standard form as:

$$\max_{y \geq 0} \{\bar{1}^T y : Ay \leq \bar{1}\}. \tag{2}$$

We denote by OPT the optimal value of the primal (1) (which is also the optimal value of the dual (2)). In this case, we say that x is a $(1 + \epsilon)$ -approximation for the covering LP if $Ax \geq \bar{1}$ and $\bar{1}^T x \leq (1 + \epsilon)$ OPT, and we say that y is a $(1 - \epsilon)$ -approximation for the packing LP if $Ay \leq \bar{1}$ and $\bar{1}^T y \geq (1 - \epsilon)$ OPT.

*Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720. Email: wangd@eecs.berkeley.edu

†Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720. Email: satishr@berkeley.edu

‡International Computer Science Institute and Department of Statistics, University of California at Berkeley, Berkeley, CA 94720. Email: mmahoney@stat.berkeley.edu

Packing and covering problems are important classes of LPs with wide applications, and they have long drawn interest in computer science and theoretical computer science. Although one can use general LP solvers such as interior point method to solve packing and covering with convergence rate of $\log(1/\epsilon)$, such algorithms usually have very high per-iteration cost, as methods such as the computation of the Hessian and matrix inversion are involved. In the setting of large-scale problems, low precision iterative solvers are often more popular choices. Such solvers usually run in time with a nearly-linear dependence on the problem size, and they have $\text{poly}(1/\epsilon)$ dependence on the approximation parameter. Most such work falls into one of two categories. The first category follows the approach of transforming LPs to convex optimization problems, then applying efficient first-order optimization algorithms. Examples of work in this category include [1–3, 7, 8, 11], and all except [1, 2] apply to more general classes of LPs. The second category is based on the Lagrangian relaxation framework, and some examples of work in this category include [4–6, 10, 12, 13]. For a more detailed comparison of this prior work, see Table 1 in [1]. Also, based on whether the running time depends on the width ρ , a parameter which typically depends on the dimension and the largest entry of A , these algorithms can also be divided into width-dependent solvers and width-independent solvers. Width-dependent solvers are usually pseudo-polynomial, as the running time depends on ρOPT , which itself can be large, while width-independent solvers are more efficient in the sense that they provide truly polynomial-time approximation solvers.

In this paper, we describe a solver for covering LPs of the form (1). The solver is width-independent,¹ and it is a first-order method with a linear rate of convergence. That is, if we let N be the number of non-zeros in A , then the running time of our algorithm is at worst $O\left(N \frac{\log^2(N/\epsilon) \log(1/\epsilon)}{\epsilon}\right)$. To simplify the following discussion, we will follow the standard practice of using \tilde{O} to hide poly-log factors, in which case the running time of our algorithm for the covering problem is at worst $\tilde{O}(N/\epsilon)$. Among other things, our result is an improvement over the recent bound of $\tilde{O}(N/\epsilon^{1.5})$ provided by Allen-Zhu and Orecchia for the covering problem using a different more complicated algorithm [1], and our result corresponds to the linear rate of convergence that accelerated gradient methods are designed to achieve [8].

At least as interesting as the $\tilde{O}(1/\epsilon^{0.5})$ improvement for covering LPs, however, is the context of this problem and the main technical contribution that we developed and exploited to achieve our improvement.

- The context for our results has to do with the linear coupling method that was introduced recently by Allen-Zhu and Orecchia [14]. This is a method for solving convex optimization problems with first order methods, and it provides a conceptually simple way to integrate a gradient descent step and mirror descent step in each iteration. In the setting of standard smooth convex optimization, the method achieves the same convergence rate as that of the accelerated gradient descent method of Nesterov [8], and indeed the former can be viewed as an insightful reinterpretation of the latter. The high-level approach of the linear coupling method is very flexible, and it has shown initial promise by providing improved algorithms for packing and covering LPs [1, 2].

The particular motivation for our work is a striking discrepancy between bounds provided for packing and covering LPs in the recent result of Allen-Zhu and Orecchia in [1]. In particular, they provide a $(1 - \epsilon)$ -approximation solver for the packing problem in $\tilde{O}(N/\epsilon)$, but they are only able to obtain $\tilde{O}(N/\epsilon^{1.5})$ for the covering problem, and for that they need to use a different more complicated algorithm. This discrepancy between results for packing and covering LPs is rare, due to the duality between them, and it leaves open the question of the exact role that the linear coupling is playing in coordinating the complementary gradient and mirror descent step of the algorithms for these dual problems.

- Our main technical contribution is a novel diameter reduction method for fractional covering LPs that helps resolve this discrepancy. Recall that the smoothness parameter, e.g., Lipschitz constant, and the diameter of the feasible region are the two most natural limiting factors for most gradient based optimization algorithms. Indeed, many applications of general first-order optimization techniques can be attributed to the existence of norms or proximal setups for the specific problems that gives both good smoothness and diameter properties. In the particular case of coordinate descent algorithms based on the linear coupling idea, we additionally need good coordinate-wise diameter properties to achieve accelerated convergence.

This is easy to accomplish for packing problems, but it is not easy to do for covering problems, and this is this difference that leads to the $\tilde{O}(1/\epsilon^{0.5})$ discrepancy between packing and covering algorithms in previous work [1]. Our diameter reduction method for general covering problems is straightforward, and it

¹More precisely, our method has a logarithmic dependence on the width, but by Observation 4.2 below, this cannot be worse than $\log(nm/\epsilon)$, and thus we consider it as width-independent.

gives both good diameter bounds with respect to the canonical norm for accelerated stochastic coordinate descent (as is needed generally [1, 9]) as well as good coordinate-wise diameter bounds (as is needed for linear coupling [1]). Thus, it is likely of interest more generally for combinatorial optimization problems.

Once the diameter reduction is achieved, the remaining work is mainly straightforward, as we can directly apply known optimization schemes that work well for problems with good diameter properties. In particular, by using the scheme from [1] that was developed for packing LPs, we obtain improved $\tilde{O}(N/\epsilon)$ results for covering LPs; and this provides a unified acceleration method (unified in the sense that it is with the same algorithm and almost identical analysis) for both packing and covering LPs.

We will start in Section 2 with a description of some of the challenges in applying acceleration techniques in a unified way to these two dual problems, including those that limited previous work. Then, in Section 3 we will present our main technical contribution, a novel diameter reduction method for any covering LP of the form given in (1). Finally, in Section 4 we describe how to combine this with previous work to obtain a unified acceleration method for packing and covering problems. We include a full description of the latter analysis, with some of the details deferred to Appendix A.

2 High-level Description of Challenges

At a high level, we (as well as Allen-Zhu and Orecchia [1, 2]) use the same two-step approach of Nesterov [8]. The first step involves smoothing, which transforms the constrained problem into a *smooth* objective function with trivial or no constraints. By smooth, we mean that the gradient of the objective function has some property in the flavor of Lipschitz continuity. Once smoothing is accomplished, the second step uses one of several first order methods for convex optimization in order to obtain an approximate solution. Examples of standard application of this approach to covering LPs includes the width-dependent solvers of [7, 8] as well as multiplicative weights update solvers [3].

The first width-independent result following the optimization approach in [2] achieves width-independence by truncating the gradient, thus effectively reducing the width to 1. The algorithm uses, in a white-box way, the coupling of mirror descent and gradient descent from [14], which can be viewed as a re-interpretation of Nesterov’s accelerated gradient method [8]. However, although [2] uses a coupling of mirror descent and gradient descent, the role of gradient descent is only for width-independence, i.e., to cover the loss incurred by the large component of the gradient (see Eqn. (7) below for the precise formulation of this loss), and it is independent of the mirror descent part acting on the truncated gradient. In addition, [2] deviates from the canonical smoothing with entropy, as it instead uses generalized entropy. Importantly, the objective function to be minimized is *not* smooth in the standard Lipschitz continuity sense, but it does satisfy a similar local Lipschitz property.

To improve the sequential packing solver in [2] with convergence $\tilde{O}(1/\epsilon^3)$ to $\tilde{O}(1/\epsilon)$, the same authors in [1] apply a stochastic coordinate descent method based on the linear coupling idea. Barring the difference between Lipschitz and local Lipschitz continuity, the results in [1] can be viewed as a variant of accelerated coordinate descent method [9]. There are two places where the algorithm achieves an improvement over prior packing-covering results.

- One factor of improvement is due to the better coordinate-wise Lipschitz constant over the full dimensional Lipschitz constant. Intuitively, in the case of packing or covering, the gradient of variable x_i depends on the penalties of constraints involving x_i , which further depend on all the variables in those constraints. As a result, if we move all the variables simultaneously, we can only take a small step before changing the gradient of x_i drastically.
- The other factor of improvement comes from accelerating the gradient method. The role of gradient descent in the packing solver of [1] is twofold. First, it covers the loss incurred by the large component of the gradient as in [2] to give width-independence. Second, to accelerate the coupling as in [14], the gradient descent also needs to cover the regret term incurred by the mirror descent step (see Eqn. (7) below for the precise formulation of this regret). The adoption of A -norm (defined in Eqn. (6) below) enables the acceleration. This A -norm works particularly well for packing problems, in the sense that it easily leads to good diameter bounds: since the packing constraints impose a naive upper bound of $x_i^* \leq 1/\|A_{:i}\|_\infty$ on each variable, thus the feasible region has a small diameter $\max_{x: f(x) \leq f(x_0)} \|x - x^*\|_A$.

The importance of the small diameter is twofold. First, the diameter naturally arises in the convergence bound of gradient based methods, so we always need to use a norm or proximal setup giving small diameter

to achieve good convergence. Second, and more importantly, in this case the small diameter $[0, 1/\|A_{:,i}\|_\infty]$ on each coordinate relates the mirror descent step length and the gradient descent step length. As the regret term in mirror descent and the improvement of gradient descent step are both proportional to their respective step lengths, the small coordinate-wise diameter makes it possible to use gradient descent improvement to cover the mirror descent regret.

The combination of gradient truncation, stochastic coordinate descent, and acceleration due to small diameter in A -norm leads to the $\tilde{O}(N/\epsilon)$ solver for the packing LP [1].

Shifting to solvers for the covering LP, one obvious obstacle to reproducing the packing result is we no longer have the small diameter in A -norm. Indeed, a naive coordinate-wise upper bound from the covering constraints only gives $x_i^* \leq 1/\min_j\{A_{ji} : A_{ji} > 0\}$. Because of this, the covering solver in [1] instead use the proximal setup in their earlier work [2]. The particular proximal setup gives a good diameter for the feasible region they use, but it doesn't give a similarly good coordinate-wise diameter to enable the acceleration. To improve upon the $O(1/\epsilon^2)$ convergence of standard mirror descent, the authors use a negative-width technique as in [3] (Theorem 3.3 with $l = \sqrt{\epsilon}$). This then leads to the (improved, but still worse than for packing) $\tilde{O}(1/\epsilon^{1.5})$ convergence rate. In addition, since they truncate the gradient at a smaller threshold to cover the loss incurred by the large component, they need a more complicated gradient step, leading to a more complicated algorithm than for the packing LP.

To get an $\tilde{O}(1/\epsilon)$ solver for the covering LP, it seems crucial to relate the gradient descent step and mirror descent step the same way as in the packing solver in [1]. Thus, we will stick with the A -norm, and we will work directly to reduce the diameter. Our main result (presented next in Section 3) is a general diameter reduction method to achieve the same diameter property as in the packing solver, and this enables us (in Section 4) to extend all the crucial ideas of the packing solver in [1], as outlined in this section, to get a covering solver with running time $\tilde{O}(N/\epsilon)$.

3 Diameter Reduction Method for General Covering Problems

Given any covering LP of the form given in (1), characterized by a matrix A , we formulate an equivalent covering LP with good diameter properties. This will involve adding variables and redundant constraints. We use $i \in [n]$ to denote the indices of the variables (i.e., columns of A) and $j \in [m]$ to denote the indices of constraints (i.e., rows of A). For ease of comparison with [1], and since our unified approach for both packing and covering uses their packing solver and a similar analysis, we use the same notation whenever possible.

For any $i \in [n]$, let

$$r_i \stackrel{\text{def}}{=} \frac{\max_j\{A_{ji} : A_{ji} > 0\}}{\min_j\{A_{ji} : A_{ji} > 0\}},$$

be the ratio between the largest non-zero coefficient and the smallest non-zero coefficient of variable x_i in all constraints, and let $n_i \stackrel{\text{def}}{=} \lceil \log r_i \rceil$. We first duplicate each original variable n_i times to obtain $\bar{x}_{(i,l)}, i \in [n], l \in [n_i]$ as the new variables. In terms of the coefficient matrix, we now have a new matrix, call it $\bar{A} \in \mathbb{R}_{\geq 0}^{m \times (\sum_i n_i)}$, which contains n_i copies of the i -th column $A_{:,i}$. We denote a column of \bar{A} by the tuple (i, l) with $l \in [n_i]$. Obviously, the covering LP given by \bar{A} is equivalent to the original covering LP given by A . Adding additional copies of variables, however, will allow us to improve the diameter. To reduce the diameter of this new covering LP, we further decrease some of the coefficients in \bar{A} , and we put upper bounds on the variables. In particular, for j, i, l , we have

$$\bar{A}_{j,(i,l)} = \min\{A_{j,i}, 2^l \min_j\{A_{ji} : A_{ji} > 0\}\}, \quad (3)$$

and for variable $\bar{x}_{(i,l)}$, we add the constraint

$$\bar{x}_{(i,l)} \leq \frac{2}{2^l \min_j\{A_{ji} : A_{ji} > 0\}}. \quad (4)$$

The next lemma shows that the covering LP given by \bar{A} and the covering LP given by A are equivalent.

Lemma 3.1. *Let OPT be the optimal value of the covering LP given by A , and let $\overline{\text{OPT}}$ be the optimal of the covering LP given by \bar{A} and (4), as constructed above; then $\text{OPT} = \overline{\text{OPT}}$.*

Proof. Given any feasible solution \bar{x} , consider the solution x where $x_i = \sum_{l=1}^{n_i} \bar{x}_{(i,l)}$. It is obvious $\bar{\Gamma}^T x = \bar{\Gamma}^T \bar{x}$, and $Ax \geq \bar{\Gamma}$, as coefficients in \bar{A} are no larger than coefficients in A . Thus $\text{OPT} \leq \overline{\text{OPT}}$.

For the other direction, consider any feasible x . For each i , we can assume without loss of generality that

$$x_i \leq \frac{1}{\min_j \{A_{ji} : A_{ji} > 0\}}.$$

Let l_i be the largest index such that

$$x_i \leq \frac{2}{2^{l_i} \min_j \{A_{ji} : A_{ji} > 0\}},$$

and then let

$$\bar{x}_{(i,l)} = \begin{cases} x_i & \text{if } l = l_i \\ 0 & \text{if } l \neq l_i \end{cases}.$$

By construction, \bar{x} satisfies all the upper bounds described in (4). Furthermore, for constraint j , we must have $\bar{A}_{j,\bar{x}} \geq 1$. Since for any i , $\bar{A}_{j,(i,l_i)}$ differs from A_{ji} only when $A_{ji} > 2^{l_i} \min_j \{A_{ji} : A_{ji} > 0\}$, and we must have $l_i < n_i$ in this case by definition of n_i , which gives $\bar{x}_{(i,l_i)} = x_i \geq \frac{1}{2^{l_i} \min_j \{A_{ji} : A_{ji} > 0\}}$ by our choice of l_i being the largest possible. Then we know $\bar{A}_{j,(i,l_i)} = 2^{l_i} \min_j \{A_{ji} : A_{ji} > 0\}$, so the j -th constraint is satisfied. Thus $\text{OPT} \geq \overline{\text{OPT}}$, and we can conclude $\text{OPT} = \overline{\text{OPT}}$. \square

Given that we have shown that the covering LP defined by \bar{A} and that defined by A are equivalent, we now point out that the seemingly-redundant constraints of (4) turn out to be crucial. The reason is that the feasible region now has a small diameter in the coordinate-wise weighted 2-norm $\|\cdot\|_A$. In particular, we can rewrite the constraints (4) to be

$$\bar{x}_{(i,l)} \leq \frac{2}{\|\bar{A}_{:, (i,l)}\|_\infty}.$$

For any i , this is the same upper bound on $\bar{x}_{(i,l)}$ for $l < n_i$ (consider the row $j^* = \arg\max_j \{A_{ji}, A_{ji} > 0\}$), and it is a relaxation on $\bar{x}_{(i,n_i)}$.

The price we pay for this diameter improvement is that the new LP defined by \bar{A} is larger than that defined by A . Two comments on this are in order. First, by Observation 4.2 below, r_i is bounded by $n^2 m / \epsilon^2$, and so the diameter reduction step only increases the problem size by $O(\log(mn/\epsilon))$. Second, we have presented our diameter reduction as an explicit pre-processing step so we can use one unified optimization algorithm (Algorithm 1 below) for both packing and covering, but in practice the diameter reduction would not have to be carried out explicitly. It can equivalently be implemented implicitly within the algorithm (a trivially-modified version of Algorithm 1 below) by randomly choosing a scale after picking the coordinate i and then computing $\bar{A}_{j,(i,l)}$ in (3) by shifting bits on the fly.

Given this reduction, in the rest of the paper, when we refer to the covering LP, we will implicitly be referring to the diameter reduced version, and we have the additional guarantee that there exists an optimal solution x^* to (1) such that

$$0 \leq x_i^* \leq \frac{2}{\|A_{:,i}\|_\infty} \quad \forall i \in [n]. \quad (5)$$

4 An Accelerated Solver for (Packing and) Covering LPs

In this section, we will present our solver for covering LPs of the form (1). To motivate this, recall that for packing problems of the form (2), bounds of the form (5) automatically follow from the packing constraints $Ax \leq \bar{\Gamma}$. For readers familiar with the packing LP solver in [1], it should be plausible that—once we have this diameter property—the same stochastic coordinate descent optimization scheme will lead to a $\tilde{O}(N/\epsilon)$ covering LP solver. We now show that indeed the same optimization algorithm for packing LPs can be easily extended to solving covering LPs, thus establishing a unified acceleration method for packing and covering problems.

In Section 4.1, we'll present some preliminaries and describe how we perform smoothing on the original covering objective function; and then in Section 4.2, we'll present our main algorithm. This algorithm involves a mirror descent step, that will be described in Section 4.3, a gradient descent step, that will be described in Section 4.4, and a careful coupling between the two, that will be described in Section 4.5. Finally, in Section 4.6,

we will describe how to ensure we start at a good starting point. Some of the following results are technically-tedious but conceptually-straightforward extensions of analogous results from [1], and some of the results are restated from [1]; for completeness, we provide the proof of all of these results, with the latter relegated to Appendix A.

4.1 Preliminaries and Smoothing the Objective

To start, let's assume that

$$\min_{j \in [m]} \|A_{j:}\|_{\infty} = 1.$$

This assumption is without loss of generality: since we are interested in multiplicative $(1 + \epsilon)$ -approximation, we can simply scale A for this to hold without sacrificing approximation quality. With this assumption, the following lemma holds. (This lemma is the same as Proposition C.2.(a) in [1], and its proof is included for completeness in Appendix A.)

Lemma 4.1. $\text{OPT} \in [1, m]$

With OPT being at least 1, the error we introduce later in the smoothing step will be small enough that the smoothing function approximates the covering LP well enough with respect to ϵ around the optimum.

Observation 4.2. *Since we are interested in a $(1 + \epsilon)$ -approximation, then with the above assumption, we can also eliminate the very small and very large entries from the matrix as follows. If some entry $A_{ji} \leq \epsilon/(mn)$, then since $\text{OPT} \leq m$ we have that $A_{ji}x_i^* \leq \epsilon/n$, and so we can just increase each variable by ϵ/n , in which case we can recover the loss from setting A_{ji} equal to 0 from the variable in the j -th constraint with coefficient at least 1. On the other hand, if some entry $A_{ji} \geq n/\epsilon$, then we can just set variable i to be at least ϵ/n and ignore constraint j . Thus, we can eliminate very small and very large entries from the matrix A , and we only incur an additional cost of ϵ , but since $\text{OPT} \geq 1$, we still obtain a $(1 + O(\epsilon))$ -approximation.*

We will turn the covering LP objective into a smoothed objective function $f_{\mu}(x)$, as used in [1, 2], and we are going to find a $(1 + \epsilon)$ -approximation of the covering LP by approximately minimizing $f_{\mu}(x)$ over the region

$$\Delta \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n : 0 \leq x_i \leq \frac{3}{\|A_{:i}\|_{\infty}}\}.$$

The function $f_{\mu}(x)$ is

$$f_{\mu}(x) \stackrel{\text{def}}{=} \vec{1}^T x + \max_{y \geq 0} \{y^T (\vec{1} - Ax) + \mu H(y)\},$$

and it is a smoothed objective in the sense that it turns the covering constraints into soft penalties, with $H(y)$ being a regularization term. Here, we use the generalized entropy $H(y) = -\sum_j y_j \log y_j + y_j$, where μ is the smoothing parameter balancing the penalty and the regularization. It is straightforward to compute the optimal y , and write $f_{\mu}(x)$ explicitly, as stated in the following lemma.

Lemma 4.3. $f_{\mu}(x) = \vec{1}^T x + \mu \sum_{j=1}^m p_j(x)$, where $p_j(x) \stackrel{\text{def}}{=} \exp(\frac{1}{\mu}(1 - (Ax)_j))$.

Optimizing $f_{\mu}(x)$ over Δ gives a good approximation to OPT , in the following sense. If we let x^* be an optimal solution satisfying (5), and $u^* \stackrel{\text{def}}{=} (1 + \epsilon/2)x^* \in \Delta$, then we have the properties in the following lemma. (This lemma is the same as Proposition C.2 in [1], and its proof is included for completeness in Appendix A.)

Lemma 4.4. *Setting the smoothing parameter $\mu = \frac{\epsilon}{4 \log(nm/\epsilon)}$, we have*

1. $f_{\mu}(u^*) \leq (1 + \epsilon) \text{OPT}$.
2. $f_{\mu}(x) \geq (1 - \epsilon) \text{OPT}$ for any $x \geq 0$.
3. For any $x \geq 0$ satisfying $f_{\mu}(x) \leq 2 \text{OPT}$, we must have $Ax \geq (1 - \epsilon)\vec{1}$.
4. If $x \geq 0$ satisfies $f_{\mu}(x) \leq (1 + O(\epsilon)) \text{OPT}$, then $\frac{1}{1 - \epsilon}x$ is a $(1 + O(\epsilon))$ -approximation to the covering LP.

5. The gradient of $f_\mu(x)$ is

$$\nabla f_\mu(x) = \vec{1} - A^T p(x) \quad \text{where} \quad p_j(x) \stackrel{\text{def}}{=} \exp\left(\frac{1}{\mu}(1 - (Ax)_j)\right),$$

$$\text{and } \nabla_i f_\mu(x) = 1 - \sum_j A_{ji} p_j(x) \in [-\infty, 1].$$

Although $f_\mu(x)$ gives a good approximation to the covering LP, we cannot simply apply the standard (accelerated) gradient descent algorithm to optimize it, as $f_\mu(x)$ doesn't have the necessary Lipschitz-smoothness property. However, $f_\mu(x)$ is *locally Lipschitz continuous*, in a sense quantified by the following lemma, and so we have a good improvement with a gradient step within certain range. (The following is a “symmetric” version² of Lemma 2.6 in [1].)

Lemma 4.5. Let $L \stackrel{\text{def}}{=} \frac{4}{\mu}$, for any $x \in \Delta$, and $i \in [n]$

1. If $\nabla_i f_\mu(x) \in (-1, 1)$, then for all $|\gamma| \leq \frac{1}{L\|A_{:i}\|_\infty}$, we have

$$|\nabla_i f_\mu(x) - \nabla_i f_\mu(x + \gamma \mathbf{e}_i)| \leq L\|A_{:i}\|_\infty |\gamma|.$$

2. If $\nabla_i f_\mu(x) \leq -1$, then for all $\gamma \leq \frac{1}{L\|A_{:i}\|_\infty}$, we have

$$\nabla_i f_\mu(x + \gamma \mathbf{e}_i) \leq (1 - \frac{L\|A_{:i}\|_\infty}{2} |\gamma|) \nabla_i f_\mu(x).$$

Proof. First, observe the following:

$$\begin{aligned} \left| \log \frac{1 - \nabla_i f_\mu(x + \gamma \mathbf{e}_i)}{1 - \nabla_i f_\mu(x)} \right| &= \left| \int_0^\gamma -\frac{\nabla_{ii} f_\mu(x + \nu \mathbf{e}_i)}{1 - \nabla_i f_\mu(x + \nu \mathbf{e}_i)} d\nu \right| = \left| \frac{1}{\mu} \int_0^\gamma \frac{\sum_j A_{ji}^2 p_j(x + \nu \mathbf{e}_i)}{\sum_j A_{ji} p_j(x + \nu \mathbf{e}_i)} d\nu \right| \\ &\leq \left| \frac{1}{\mu} \int_0^\gamma \|A_{:i}\|_\infty d\nu \right| = \frac{1}{\mu} |\gamma| \|A_{:i}\|_\infty = \frac{L\|A_{:i}\|_\infty}{4} |\gamma|. \end{aligned}$$

Then, we have

$$\exp\left(-\frac{L\|A_{:i}\|_\infty}{4} |\gamma|\right) \leq \frac{1 - \nabla_i f_\mu(x + \gamma \mathbf{e}_i)}{1 - \nabla_i f_\mu(x)} \leq \exp\left(\frac{L\|A_{:i}\|_\infty}{4} |\gamma|\right).$$

Since $\frac{L\|A_{:i}\|_\infty}{4} |\gamma| \leq \frac{1}{4}$ by our assumption, we have $x \leq e^x - 1 \leq 1.2x$ for $x \in [-\frac{1}{4}, \frac{1}{4}]$. Thus, it follows that

$$-\frac{L\|A_{:i}\|_\infty}{4} |\gamma| \leq \frac{\nabla_i f_\mu(x) - \nabla_i f_\mu(x + \gamma \mathbf{e}_i)}{1 - \nabla_i f_\mu(x)} \leq 1.2 \frac{L\|A_{:i}\|_\infty}{4} |\gamma|.$$

Finally, to prove the lemma we consider the following two cases:

1. If $\nabla_i f_\mu(x) \in (-1, 1)$, then we have

$$|\nabla_i f_\mu(x) - \nabla_i f_\mu(x + \gamma \mathbf{e}_i)| \leq 1.2(1 - \nabla_i f_\mu(x)) \frac{L\|A_{:i}\|_\infty}{4} |\gamma| \leq L\|A_{:i}\|_\infty |\gamma|.$$

2. If $\nabla_i f_\mu(x) \leq -1$, then $1 - \nabla_i f_\mu(x) \leq -2\nabla_i f_\mu(x)$, and

$$\nabla_i f_\mu(x + \gamma \mathbf{e}_i) \leq \nabla_i f_\mu(x) + (1 - \nabla_i f_\mu(x)) \frac{L\|A_{:i}\|_\infty}{4} |\gamma| \leq (1 - \frac{L\|A_{:i}\|_\infty}{2} |\gamma|) \nabla_i f_\mu(x).$$

□

We call $L\|A_{:i}\|_\infty$ the *coordinate-wise local Lipschitz constant*. For readers familiar with accelerated coordinate descent method (ACDM) [9], the A -norm is essentially the $\|\cdot\|_{1-\alpha}$ in ACDM [9] with $\alpha = 0$, except we use the coordinate-wise local Lipschitz constant instead of the Lipschitz constant to weight each coordinate. The significance of Lemma 4.5 is that for covering LPs the coordinate-wise diameter is inversely proportional to the coordinate-wise local Lipschitz constant. (This fact has been established previously for the case of packing LPs [1].)

Algorithm 1 Accelerated stochastic coordinate descent for both packing and covering

Input: $A \in \mathbb{R}_{\geq 0}^{m \times n}$, $x^{\text{start}} \in \Delta$, f_μ, ϵ **Output:** $y_T \in \Delta$

- 1: $\mu \leftarrow \frac{\epsilon}{4 \log(nm/\epsilon)}, L \leftarrow \frac{4}{\mu}, \tau \leftarrow \frac{1}{8nL}$
 - 2: $T \leftarrow \lceil 8nL \log(1/\epsilon) \rceil = \tilde{O}(\frac{n}{\epsilon})$
 - 3: $x_0, y_0, z_0 \leftarrow x^{\text{start}}, \alpha_0 \leftarrow \frac{1}{nL}$
 - 4: **for** $k = 1$ to T **do**
 - 5: $\alpha_k \leftarrow \frac{1}{1-\tau} \alpha_{k-1}$
 - 6: $x_k \leftarrow \tau z_{k-1} + (1-\tau)y_{k-1}$
 - 7: Select $i \in [n]$ uniformly at random.
 - 8: \triangleright Gradient truncation:
Let $\xi_k^{(i)} \leftarrow \begin{cases} -1 & \nabla_i f_\mu(x_k) < -1 \\ \nabla_i f_\mu(x_k) & \nabla_i f_\mu(x_k) \in [-1, 1] \\ 1 & \nabla_i f_\mu(x_k) > 1 \end{cases}$
 - 9: \triangleright Mirror descent step:
 $z_k \leftarrow z_k^{(i)} \stackrel{\text{def}}{=} \operatorname{argmin}_{z \in \Delta} \{V_{z_{k-1}}(z) + \langle z, n\alpha_k \xi_k^{(i)} \rangle\}.$
 - 10: \triangleright Gradient descent step:
 $y_k \leftarrow y_k^{(i)} \stackrel{\text{def}}{=} x_k + \frac{1}{n\alpha_k L} (z_k^{(i)} - z_{k-1})$
 - 11: **end for**
 - 12: **return** y_T .
-

4.2 An Accelerated Coordinate Descent Algorithm

We will now show that the accelerated coordinate descent used in packing LP solver in [1] also works as a covering LP solver, with appropriately-chosen starting points and smoothed objective functions. Consider Algorithm 1, which is our main accelerated stochastic coordinate descent for both packing and covering. This algorithm takes as input a matrix $A \in \mathbb{R}_{\geq 0}^{m \times n}$, an initial condition $x^{\text{start}} \in \Delta$, a smoothed function f_μ , and an error parameter ϵ , and it returns as output a vector $y_T \in \Delta$. The correctness of this algorithm and its running time guarantees for the packing problem have already been nicely presented in [1], and so here we will focus on the covering problem.

Our main result is summarized in the following theorems.

Theorem 4.6. *With x^{start} computable in time $\tilde{O}(N)$ to be specified later, Algorithm 1 outputs y_T satisfying $\mathbb{E}[f_\mu(y_T)] \leq (1 + 6\epsilon) \text{OPT}$, and the running time is $\tilde{O}(N/\epsilon)$.*

Given Theorem 4.6, a standard application of Markov bound, together with part 5 of Lemma 4.4, gives the following theorem as a corollary.

Theorem 4.7. *There is a algorithm that, with probability at least 9/10, computes a $(1 + O(\epsilon))$ -approximation to the fractional covering problem and has $\tilde{O}(N/\epsilon)$ expected running time.*

Not surprisingly, due to the structural similarities of packing and covering problems after diameter reduction, the correctness of Algorithm 1 for covering can be established using the same approach as [1] did for packing. The modifications are fairly straightforward, and we will point out the similarities whenever possible.

Before proceeding with our proof of these theorems, we discuss briefly the optimization scheme from [1] we will use. First, observe that the A -norm, where

$$\|x\|_A = \sqrt{\sum_i \|A_{:i}\|_\infty x_i^2}, \quad (6)$$

is used as the proximal setup for mirror descent. The corresponding distance generating function is $w(x) = \frac{1}{2} \|x\|_A^2$, and the Bregman divergence is $V_x(y) = \frac{1}{2} \|x - y\|_A^2$.³

²The smoothed objective function for packing LP is $-\bar{1}^T y + \mu \sum_{j=1}^m q_j(y)$, where $q_j(y) \stackrel{\text{def}}{=} \exp(\frac{1}{\mu}((Ay)_j - 1))$, which is symmetric to $f_\mu(x)$. The properties of $f_\mu(x)$ inherit the symmetry to its packing counterpart, and it can be derived with the same way as [1] used for the packing function, but we include it's proof to highlight differences.

³In particular, w is a 1-strongly convex function with respect to $\|\cdot\|_A$, and $V_x(y) \stackrel{\text{def}}{=} w(y) - \langle \nabla w(x), y - x \rangle - w(x)$. See [14] for a detailed discussion of mirror descent as well as and several interpretations.

Next, observe that Algorithm 1 works as follows. Each iteration integrates a mirror descent step and a gradient descent step. The standard analysis of mirror descent gives a convergence of $\frac{1}{\epsilon^2}$, and it depends on the width of the problem. Thus, to get a width-independent $\tilde{O}(\frac{N}{\epsilon})$ solver, we need to show that Algorithm 1 addresses both of these issues.

- In order to eliminate the width from the convergence rate, the gradient $\nabla_i f_\mu(x_k)$ is split into the small component, $\xi_k^{(i)} = \max\{-1, \nabla_i f_\mu(x_k)\} \mathbf{e}_i$, and the large component, $\eta_k^{(i)} = \nabla_i f_\mu(x_k) \mathbf{e}_i - \xi_k^{(i)}$. Only the small component $\xi_k^{(i)}$ is given to the mirror descent step, and thus the width is effectively 1. However, the truncation incurs loss from the large component, as the mirror descent only acts on the small component. Following [2], the improvement from the gradient descent step is used to cover that loss.
- In order to improve the $1/\epsilon^2$ rate, recall that the $1/\epsilon^2$ in the convergence of mirror descent is largely due to the regret term accumulated along all iterations of mirror descent. In order to get to $1/\epsilon$, the improvement from the gradient step also need to cover the regret from the mirror descent step (see Eqn. (7) below for the precise formulation of this loss and regret). This enables us to telescope both the loss and the regret through all iterations and to bound the total by the gap between $f_\mu(x^{\text{start}})$ and the optimal. The remaining terms in the mirror descent also telescope through the algorithm, and they are bounded in total by the distance (in A -norm) from x^{start} to $u^* \in \Delta$.

Then, given these, all we need is an initial condition x^{start} that is not too far away from the optimal in terms of the function value and not too far away from u^* in A -norm. For packing, starting with all 0's will work. For covering, we will show later a good enough x^{start} can be obtained in $\tilde{O}(N)$.

Finally, here are some lemmas about the algorithm. The following two lemmas are invariant to the differences between packing and covering problems, and so they follow directly from the same results in [1] (but, for completeness, we include the proofs in Appendix A). The values of parameters μ, L, τ, α_k can be found in the description of Algorithm 1. The first lemma says that the gradient step we take is always valid (i.e., in Δ), which is crucial in the sense that the gradient descent improvement is proportional to the step length, and we need the step length to be at least $\frac{1}{n\alpha_k L}$ of the mirror descent step length for the coupling to work.

Lemma 4.8. *We have $x_k, y_k, z_k \in \Delta$ for all $k = 0, 1, \dots, T$.*

The second lemma is clearly crucial to achieve the nearly linear time $\tilde{O}(N/\epsilon)$ algorithm.

Lemma 4.9. *Each iteration can be implemented in expected $O(N/n)$ time.*

4.3 Mirror Descent Step

We now analyze the mirror descent step of Algorithm 1:

$$z_k \leftarrow z_k^{(i)} \stackrel{\text{def}}{=} \underset{z \in \Delta}{\operatorname{argmin}} \{V_{z_{k-1}}(z) + \langle z, n\alpha_k \xi_k^{(i)} \rangle\}.$$

A lemma of the following form, which here applies to both covering and packing LPs, is needed, and it's proof follows from the textbook mirror descent analysis (or, e.g., Lemma 3.5 in [1]).

Lemma 4.10. $\langle n\alpha_k \xi_k^{(i)}, z_{k-1} - u^* \rangle \leq n^2 \alpha_k^2 L \langle \xi^{(i)}, x_k - y_k^{(i)} \rangle + V_{z_{k-1}}(u^*) - V_{z_k}(u^*)$

Proof. The lemma follows from the following chain of equalities and inequalities.

$$\begin{aligned} \langle n\alpha_k \xi_k^{(i)}, z_{k-1} - u^* \rangle &= \langle n\alpha_k \xi_k^{(i)}, z_{k-1} - z_k \rangle + \langle n\alpha_k \xi_k^{(i)}, z_k - u^* \rangle \\ &= n^2 \alpha_k^2 L \langle \xi^{(i)}, x_k - y_k^{(i)} \rangle + \langle n\alpha_k \xi_k^{(i)}, z_k - u^* \rangle \\ &\leq n^2 \alpha_k^2 L \langle \xi^{(i)}, x_k - y_k^{(i)} \rangle + \langle -\nabla V_{z_{k-1}}(z_k^{(i)}), z_k - u^* \rangle \\ &\leq n^2 \alpha_k^2 L \langle \xi^{(i)}, x_k - y_k^{(i)} \rangle + V_{z_{k-1}}(u^*) - V_{z_k^{(i)}}(u^*) - V_{z_{k-1}}(z_k^{(i)}) \\ &\leq n^2 \alpha_k^2 L \langle \xi^{(i)}, x_k - y_k^{(i)} \rangle + V_{z_{k-1}}(u^*) - V_{z_k}(u^*). \end{aligned}$$

The first equality follows by adding and subtracting z_k , and the second equality comes from the gradient step $y_k^{(i)} = x_k + \frac{1}{n\alpha_k L}(z_k^{(i)} - z_{k-1})$. The first inequality is due to the minimality of $z_k^{(i)}$, which gives

$$\langle \nabla V_{z_{k-1}}(z_k^{(i)}) + n\alpha_k \xi_k^{(i)}, u - z_k \rangle \geq 0 \quad \forall u \in \Delta,$$

the second inequality is due to the standard three point property of Bregman divergence, that is $\forall x, y \geq 0$

$$\langle -\nabla V_x(y), y - u \rangle = V_x(u) - V_y(u) - V_x(y),$$

and the last inequality just drops the term $-V_{z_k}(u^*)$, which is always negative. \square

Also, we note that the mirror descent step, defined above in a variational way, can be explicitly written as

1. $z_k^{(i)} \leftarrow z_{k-1}$
2. $z_k^{(i)} \leftarrow z_k^{(i)} - n\alpha_k \xi_k^{(i)} / \|A_{:,i}\|_\infty$
3. If $z_{k,i}^{(i)} < 0$, $z_{k,i}^{(i)} \leftarrow 0$; if $z_{k,i}^{(i)} > 3/\|A_{:,i}\|_\infty$, $z_{k,i}^{(i)} \leftarrow 3/\|A_{:,i}\|_\infty$.

This is invariant to the difference of packing and covering, and so it follows directly from Proposition 3.6 in [1]. It is fairly easy to derive, and so we omit the proof.

4.4 Gradient Descent Step

We now analyze the gradient descent step of Algorithm 1. In particular, from the explicit formulation of the mirror descent step, we have that $|z_{k,i}^{(i)} - z_{k-1,i}| \leq \frac{n\alpha_k |\xi_k^{(i)}|}{\|A_{:,i}\|_\infty}$, which gives

$$|y_{k,i}^{(i)} - x_{k,i}| = \frac{1}{n\alpha_k L} |z_{k,i}^{(i)} - z_{k-1,i}| \leq \frac{|\xi_k^{(i)}|}{L\|A_{:,i}\|_\infty}.$$

The gradient step we take is within the local region, and so Lemma 4.5 applies. We bound the improvement from the gradient descent step in the following lemma, which is symmetric⁴ to Lemma 3.8 in [1].

Lemma 4.11. $f_\mu(x_k) - f_\mu(y_k^{(i)}) \geq \frac{1}{2} \langle \nabla f_\mu(x_k), x_k - y_k^{(i)} \rangle$

Proof. Since x_k and $y_k^{(i)}$ differ only at coordinate i , denote $\gamma = y_{k,i}^{(i)} - x_{k,i}$, we have

$$f_\mu(x_k) - f_\mu(y_k^{(i)}) = f_\mu(x_k) - f_\mu(x_k + \gamma \mathbf{e}_i) = \int_0^\gamma -\nabla_i f_\mu(x_k + \nu \mathbf{e}_i) d\nu.$$

Since γ satisfies $|\gamma| \leq \frac{|\xi_k^{(i)}|}{L\|A_{:,i}\|_\infty} \leq \frac{1}{L\|A_{:,i}\|_\infty}$, we can apply Lemma 4.5. There are two cases to consider.

If $\nabla_i f_\mu(x_k) \in (-1, 1)$, then we have $|\gamma| \leq \frac{|\xi_k^{(i)}|}{L\|A_{:,i}\|_\infty} = \frac{|\nabla_i f_\mu(x_k)|}{L\|A_{:,i}\|_\infty}$, and by Lemma 4.5 we have $-\nabla_i f_\mu(x_k + \nu \mathbf{e}_i) \geq -\nabla_i f_\mu(x_k) - L\|A_{:,i}\|_\infty |\nu|$ in the above integration. Thus,

$$\begin{aligned} f_\mu(x_k) - f_\mu(y_k^{(i)}) &\geq \int_0^\gamma -\nabla_i f_\mu(x_k + \nu \mathbf{e}_i) d\nu \\ &\geq \int_0^\gamma -\nabla_i f_\mu(x_k) - L\|A_{:,i}\|_\infty |\nu| d\nu \\ &= -\nabla_i f_\mu(x_k) \gamma - \frac{L\|A_{:,i}\|_\infty}{2} \gamma^2 \\ &\geq -\nabla_i f_\mu(x_k) \gamma - \frac{L\|A_{:,i}\|_\infty}{2} |\gamma| \frac{|\nabla_i f_\mu(x_k)|}{L\|A_{:,i}\|_\infty} \\ &= -\frac{1}{2} \langle \nabla_i f_\mu(x_k), \gamma \rangle = \frac{1}{2} \langle \nabla f_\mu(x_k), x_k - y_k^{(i)} \rangle. \end{aligned}$$

If $\nabla_i f_\mu(x_k) \leq -1$, then again by Lemma 4.5 we have $-\nabla_i f_\mu(x_k + \nu \mathbf{e}_i) \geq -(1 - \frac{L\|A_{:,i}\|_\infty}{2} |\nu|) \nabla_i f_\mu(x_k) \geq -\frac{1}{2} \nabla_i f_\mu(x_k)$. Thus,

$$\begin{aligned} f_\mu(x_k) - f_\mu(y_k^{(i)}) &\geq \int_0^\gamma -\nabla_i f_\mu(x_k + \nu \mathbf{e}_i) d\nu \\ &\geq \int_0^\gamma -\frac{1}{2} \nabla_i f_\mu(x_k) d\nu = \frac{1}{2} \langle \nabla f_\mu(x_k), x_k - y_k^{(i)} \rangle. \end{aligned}$$

\square

⁴The symmetry is between Lemma 2.6 in [1] and Lemma 4.5, as the gradient descent improvement follows directly from the corresponding Lipschitz properties. The actual improvement guarantee is the same as Lemma 3.8 in [1].

4.5 Coupling of Gradient and Mirror Descent

Here, we will analyze the coupling between the gradient descent and mirror descent steps. This and the next section will give a proof of Theorem 4.6.

As we take steps on random coordinates, we will write the full gradient as

$$\nabla f_\mu(x_k) = \mathbb{E}_i[n\nabla_i f_\mu(x_k)] = \mathbb{E}_i[n\eta_k^{(i)} + n\xi_k^{(i)}].$$

As discussed earlier, we have the small component $\xi_k^{(i)} \in (-1, 1) \mathbf{e}_i$ and the large component $\eta_k^{(i)} = \nabla_i f_\mu(x_k) - \xi_k^{(i)} \in (-\infty, 0] \mathbf{e}_i$. We put the gradient and mirror descent steps together, and we bound the gap to optimality at iteration k :

$$\begin{aligned} \alpha_k(f_\mu(x_k) - f_\mu(u^*)) &\leq \langle \alpha_k \nabla f_\mu(x_k), x_k - u^* \rangle \\ &= \langle \alpha_k \nabla f_\mu(x_k), x_k - z_{k-1} \rangle + \langle \alpha_k \nabla f_\mu(x_k), z_{k-1} - u^* \rangle \\ &= \langle \alpha_k \nabla f_\mu(x_k), x_k - z_{k-1} \rangle + \mathbb{E}_i[\langle n\alpha_k \eta_k^{(i)}, z_{k-1} - u^* \rangle + \langle n\alpha_k \xi_k^{(i)}, z_{k-1} - u^* \rangle] \\ &= \frac{1-\tau}{\tau} \alpha_k \langle \nabla f_\mu(x_k), y_{k-1} - x_k \rangle + \mathbb{E}_i[\langle n\alpha_k \eta_k^{(i)}, z_{k-1} - u^* \rangle] \\ &\quad + \mathbb{E}_i[\langle n\alpha_k \xi_k^{(i)}, z_{k-1} - u^* \rangle] \\ &\leq \frac{1-\tau}{\tau} \alpha_k (f_\mu(y_{k-1}) - f_\mu(x_k)) + \mathbb{E}_i[\langle n\alpha_k \eta_k^{(i)}, z_{k-1} - u^* \rangle] \\ &\quad + \mathbb{E}_i[n^2 \alpha_k^2 L \langle \xi_k^{(i)}, x_k - y_k^{(i)} \rangle + V_{z_{k-1}}(u^*) - V_{z_k^{(i)}}(u^*)]. \end{aligned}$$

The first line is due to convexity. The next two lines just break and regroup the terms. The fourth line is due to $x_k = \tau z_{k-1} + (1-\tau)y_{k-1}$, so $\tau(x_k - z_{k-1}) = (1-\tau)(y_{k-1} - x_k)$. The last line is by Lemma 4.10.

We try to use the improvement from the gradient step given in Lemma 4.11 to cover the loss from $\eta_k^{(i)}$, and the regret from the mirror descent step:

$$\underbrace{\mathbb{E}_i[\langle n\alpha_k \eta_k^{(i)}, z_{k-1} - u^* \rangle]}_{\text{loss from } \eta_k^{(i)}} + \underbrace{\mathbb{E}_i[n^2 \alpha_k^2 L \langle \xi_k^{(i)}, x_k - y_k^{(i)} \rangle]}_{\text{regret from mirror descent}}, \quad (7)$$

and we will use the fact $z_{k-1}, z_k^{(i)}, u^* \in \Delta$. Consider the following cases.

1. $\eta_k^{(i)} = 0$: In this case, the loss term is 0. We only need to worry about the regret term, and by Lemma 4.11

$$n^2 \alpha_k^2 L \langle \xi_k^{(i)}, x_k - y_k^{(i)} \rangle \leq 2n^2 \alpha_k^2 L (f_\mu(x_k) - f_\mu(y_k^{(i)})).$$

2. $\eta_k^{(i)} < 0, z_{k,i}^{(i)} < \frac{3}{\|A_{:,i}\|_\infty}$: In this case, we increased the i -th variable in both the gradient and mirror descent step, and because $z_{k,i}^{(i)}$ is inside Δ without any projection, we know the step length of gradient descent is exactly $y_{k,i}^{(i)} - x_{k,i} = \frac{1}{n\alpha_k L \|A_{:,i}\|_\infty} = \frac{1}{L\|A_{:,i}\|_\infty}$, together with $z_{k-1} \geq 0$, and $u_i^* \leq \frac{3}{\|A_{:,i}\|_\infty}$, we have

$$\langle n\alpha_k \eta_k^{(i)}, z_{k-1} - u^* \rangle \leq \langle n\alpha_k \eta_k^{(i)}, -u^* \rangle \leq -n\alpha_k \nabla_i f_\mu(x_k) \frac{3}{\|A_{:,i}\|_\infty} = 3n\alpha_k L \langle \nabla f_\mu(x_k), x_k - y_k^{(i)} \rangle,$$

and

$$\begin{aligned} \langle n\alpha_k \eta_k^{(i)}, z_{k-1} - u^* \rangle + n^2 \alpha_k^2 L \langle \xi_k^{(i)}, x_k - y_k^{(i)} \rangle &\leq (3n\alpha_k L + n^2 \alpha_k^2 L) \langle \nabla f_\mu(x_k), x_k - y_k^{(i)} \rangle \\ &\leq (6n\alpha_k L + 2n^2 \alpha_k^2 L) (f_\mu(x_k) - f_\mu(y_k^{(i)})). \end{aligned}$$

The last step is by Lemma 4.11.

3. $\eta_k^{(i)} < 0, z_{k,i}^{(i)} = \frac{3}{\|A_{:,i}\|_\infty}$: In this case, as we know $u_i^* \leq \frac{3}{\|A_{:,i}\|_\infty}$, we have

$$\langle n\alpha_k \eta_k^{(i)}, z_{k-1} - u^* \rangle \leq \langle n\alpha_k \eta_k^{(i)}, z_{k-1} - z_k^{(i)} \rangle = n^2 \alpha_k^2 L \langle \eta_k^{(i)}, x_k - y_k^{(i)} \rangle,$$

and

$$\begin{aligned} \langle n\alpha_k \eta_k^{(i)}, z_{k-1} - u^* \rangle + n^2 \alpha_k^2 L \langle \xi_k^{(i)}, x_k - y_k^{(i)} \rangle &\leq 2n^2 \alpha_k^2 L \langle \nabla f_\mu(x_k), x_k - y_k^{(i)} \rangle \\ &\leq 4n^2 \alpha_k^2 L (f_\mu(x_k) - f_\mu(y_k^{(i)})). \end{aligned}$$

Again, the last step is due to Lemma 4.11.

Since $n\alpha_k < 1$ for all k , we have in all above cases,

$$\mathbb{E}_i[\langle n\alpha_k \eta_k^{(i)}, z_{k-1} - u^* \rangle] + \mathbb{E}_i[n^2 \alpha_k^2 L \langle \xi_k^{(i)}, x_k - y_k^{(i)} \rangle] \leq \mathbb{E}_i[8n\alpha_k L (f_\mu(x_k) - f_\mu(y_k^{(i)}))].$$

Back to our earlier derivation, we have

$$\begin{aligned} \alpha_k (f_\mu(x_k) - f_\mu(u^*)) &\leq \frac{1-\tau}{\tau} \alpha_k (f_\mu(y_{k-1}) - f_\mu(x_k)) + \mathbb{E}_i[\langle n\alpha_k \eta_k^{(i)}, z_{k-1} - u^* \rangle] \\ &\quad + \mathbb{E}_i[n^2 \alpha_k^2 L \langle \xi_k^{(i)}, x_k - y_k^{(i)} \rangle + V_{z_{k-1}}(u^*) - V_{z_k^{(i)}}(u^*)] \\ &\leq \frac{1-\tau}{\tau} \alpha_k (f_\mu(y_{k-1}) - f_\mu(x_k)) + \mathbb{E}_i[8n\alpha_k L (f_\mu(x_k) - f_\mu(y_k^{(i)}))] \\ &\quad + \mathbb{E}_i[V_{z_{k-1}}(u^*) - V_{z_k^{(i)}}(u^*)]. \end{aligned}$$

With our choice of $\tau = \frac{1}{8nL}$, $\alpha_k = \frac{1}{1-\tau} \alpha_{k-1}$, we have

$$-\alpha_k f_\mu(u^*) \leq 8nL\alpha_{k-1} f_\mu(y_{k-1}) - \mathbb{E}_i[8nL\alpha_k f_\mu(y_k^{(i)})] + \mathbb{E}_i[V_{z_{k-1}}(u^*) - V_{z_k^{(i)}}(u^*)].$$

Telescoping the above inequality along $k = 1, \dots, T$, we get

$$\mathbb{E}_i[8nL\alpha_T f_\mu(y_T)] \leq \sum_{k=1}^T \alpha_k f_\mu(u^*) + 8nL\alpha_0 f_\mu(y_0) + V_{z_0}(u^*),$$

and thus

$$\mathbb{E}_i[f_\mu(y_T)] \leq \frac{\sum_{k=1}^T \alpha_k}{8nL\alpha_T} f_\mu(u^*) + \frac{\alpha_0}{\alpha_T} f_\mu(y_0) + \frac{1}{8nL\alpha_T} V_{z_0}(u^*).$$

We have $\sum_{k=1}^T \alpha_k = \alpha_T \sum_{k=0}^{T-1} (1 - \frac{1}{8nL})^k = 8nL\alpha_T (1 - (1 - \frac{1}{8nL})^T) \leq 8nL\alpha_T$, and by our choice of $T = \lceil 8nL \log(1/\epsilon) \rceil$, we also have

$$\frac{\alpha_0}{\alpha_T} = (1 - \frac{1}{8nL})^T \leq \epsilon, \frac{1}{8nL\alpha_T} \leq \frac{\epsilon}{8nL\alpha_0} = \frac{\epsilon}{8},$$

and thus

$$\mathbb{E}_i[f_\mu(y_T)] \leq f_\mu(u^*) + \epsilon f_\mu(y_0) + \frac{\epsilon}{8} V_{z_0}(u^*).$$

4.6 Finding a Good Starting Point

Here, we will describe how to find a good starting point for the algorithm. This will permit us to establish the quality-of-approximation and running time guarantees of Theorem 4.6.

A good starting point $y_0 = x^{\text{start}}$ for Algorithm 1 is an initial condition x^{start} that is not too far away from the optimal in terms of the function value (i.e. small $f_\mu(y_0)$), and not too far away from u^* in A -norm (i.e. small $V_{z_0}(u^*)$). For packing problems, starting with all the all-0's vector will work, but this will not work for covering problems. Instead, for covering problems, we will show now a good enough x^{start} can be obtained in $\tilde{O}(N)$.

To do so, recall that we can get a 2-approximation $x^\#$ to the original covering LP in time $\tilde{O}(N)$ using various nearly linear time covering solvers, e.g., those of [5, 13]. Without loss of generality, we can assume $x_i^\# \in [0, \frac{2}{\|A_{:,i}\|_\infty}]$, since we can use the diameter reduction process as specified in Lemma 3.1 to get a equivalent solution satisfying the conditions. Then, we have the following lemma.

Lemma 4.12. *Let $x^{\text{start}} = (1 + \epsilon/2)x^\#$, we have $x^{\text{start}} \in \Delta$, $f_\mu(x^{\text{start}}) \leq 4 \text{OPT}$, and $V_{x^{\text{start}}}(u^*) \leq 6 \text{OPT}$*

Proof. It is obvious that $x^{\text{start}} \in \Delta$. Thus,

$$\vec{1}^T x^{\text{start}} = (1 + \epsilon/2) \vec{1}^T x^\# \leq (1 + \epsilon/2) 2 \text{OPT} \leq 3 \text{OPT}.$$

Furthermore, we have $Ax^{\text{start}} - \vec{1} \geq (1 + \epsilon/2)Ax^\# - \vec{1} \geq \frac{\epsilon}{2}\vec{1}$, and so

$$f_\mu(x^{\text{start}}) = \mu \sum_j p_j(x^{\text{start}}) + \vec{1}^T x^{\text{start}} \leq \mu \sum_j \exp(-\frac{\epsilon/2}{\mu}) + 3 \text{OPT} \leq \frac{\mu m}{(nm)^2} + 3 \text{OPT} < 4 \text{OPT}.$$

For the divergence, we have that

$$\begin{aligned} V_{x^{\text{start}}}(u^*) &= \frac{1}{2} \sum_i \|A_{:,i}\|_\infty (x_i^{\text{start}} - u_i^*)^2 \\ &= \frac{1}{2} \sum_i \|A_{:,i}\|_\infty ((x_i^{\text{start}})^2 + (u_i^*)^2 - 2x_i^{\text{start}}u_i^*) \\ &\leq \frac{3}{2} \sum_i x_i^{\text{start}} + u_i^* \\ &\leq \frac{3}{2} (3 \text{OPT} + \text{OPT}) \leq 6 \text{OPT}, \end{aligned}$$

which proves the lemma. \square

It is now clear that we have

$$\mathbb{E}_i[f_\mu(y_T)] \leq f_\mu(u^*) + \epsilon f_\mu(y_0) + \frac{\epsilon}{8} V_{z_0}(u^*) \leq (1 + \epsilon) \text{OPT} + 4\epsilon \text{OPT} + \epsilon \text{OPT} = (1 + 6\epsilon) \text{OPT}.$$

Thus, we have the approximation guarantee in Theorem 4.6. The running time follows directly from Lemma 4.9 and $T = \tilde{O}(n/\epsilon)$.

Acknowledgments. DW was supported by ARO Grant W911NF-12-1-0541, SR was funded by NSF Grant CCF-1118083, and MM acknowledges the support of the NSF, AFOSR, and DARPA.

Appendix A Missing Proofs

The following proofs can be found in [1], and we include them here for completeness.

Lemma 4.1. $\text{OPT} \in [1, m]$

Proof. By the assumption $\min_{j \in [m]} \|A_{j,:}\|_\infty = 1$, we know at least one constraint has all coefficients at most 1, so to satisfy that constraint, we must have the sum of the variables to be at least 1. On the other hand, since each constraint has a variable with coefficient at least 1 in it, $x = \vec{1}$ clearly satisfies all constraints, so $\text{OPT} \leq m$. \square

Lemma 4.4. Setting the smoothing parameter $\mu = \frac{\epsilon}{4 \log(nm/\epsilon)}$, we have

1. $f_\mu(u^*) \leq (1 + \epsilon) \text{OPT}$.
2. $f_\mu(x) \geq (1 - \epsilon) \text{OPT}$ for any $x \geq 0$.
3. For any $x \geq 0$ satisfying $f_\mu(x) \leq 2 \text{OPT}$, we must have $Ax \geq (1 - \epsilon)\vec{1}$.
4. If $x \geq 0$ satisfies $f_\mu(x) \leq (1 + O(\epsilon)) \text{OPT}$, then $\frac{1}{1-\epsilon}x$ is a $(1 + O(\epsilon))$ -approximation to the covering LP.
5. The gradient of $f_\mu(x)$ is

$$\nabla f_\mu(x) = \vec{1} - A^T p(\vec{x}) \quad \text{where} \quad p_j(x) \stackrel{\text{def}}{=} \exp\left(\frac{1}{\mu}(1 - (Ax)_j)\right),$$

$$\text{and } \nabla_i f_\mu(x) = 1 - \sum_j A_{ji} p_j(x) \in [-\infty, 1].$$

- Proof.* 1. Since $Ax^* \geq \vec{1}$, and $u^* = (1 + \epsilon/2)x^*$, we have $(Au^*)_j - 1 \geq \epsilon/2$ for all j . Then $p_j(u^*) \leq \exp(-\frac{1}{\mu}\frac{\epsilon}{2}) = (\frac{\epsilon}{mn})^2$, and $f_\mu(u^*) = \vec{1}^T u^* + \mu \sum_{j=1}^m p_j(u^*) \leq (1 + \epsilon/2) \text{OPT} + \mu m (\frac{\epsilon}{mn})^2 \leq (1 + \epsilon) \text{OPT}$.
2. By contradiction, suppose $f_\mu(x) < (1 - \epsilon) \text{OPT}$, since $f_\mu(x) < \text{OPT} \leq m$, we must have $p_j(x) < m/\mu$ for any j , which implies $(Ax)_j \geq 1 - \epsilon$. By definition of OPT , we have $\vec{1}^T x \geq (1 - \epsilon) \text{OPT}$, since $Ax \geq (1 - \epsilon)\vec{1}$. This gives a contradiction as $f_\mu(x) > \vec{1}^T x \geq (1 - \epsilon) \text{OPT}$.
3. By contradiction, suppose there is some j such that $(Ax)_j - 1 \leq -\epsilon$, then as in the last part, we have $\mu p_j(x) \geq \mu (\frac{mn}{\epsilon})^4 > 2 \text{OPT}$, contradicting $f_\mu(x) \leq 2 \text{OPT}$.
4. For any x satisfying $f_\mu(x) \leq (1 + O(\epsilon)) \text{OPT} \leq 2 \text{OPT}$, by last part we know $Ax \geq (1 - \epsilon)\vec{1}$, so $A(\frac{1}{1-\epsilon}x) \geq \vec{1}$. We also have $\vec{1}^T (\frac{1}{1-\epsilon}x) = \frac{1}{1-\epsilon} \vec{1}^T x < \frac{1}{1-\epsilon} f_\mu(x) \leq (1 + O(\epsilon)) \text{OPT}$.
5. This is by straightforward computation. \square

Lemma 4.8. We have $x_k, y_k, z_k \in \Delta$ for all $k = 0, 1, \dots, T$.

Proof. At the start $x_0 = y_0 = z_0 = x^{\text{start}} \in \Delta$ by assumption. z_k is always in Δ as we take the projection in the mirror descent step. If we can further show $y_k \in \Delta$ for all k , we are done, since x_k is a convex combination of y_{k-1}, z_{k-1} . To show $y_k \in \Delta$, we write y_k as a convex combination of z_0, \dots, z_k , $y_k = \sum_{l=0}^k c_k^l z_l$. At $k = 0$, we have $y_0 = z_0$, and at $k = 1$, $y_1 = x_1 + \frac{1}{n\alpha_1 L}(z_1 - z_0) = \frac{1}{n\alpha_1 L}z_1 + (1 - \frac{1}{n\alpha_1 L})z_0$, as $x_1 = y_0 = z_0$. For $k \geq 2$, we can verify

$$c_k^l = \begin{cases} (1 - \tau)c_{k-1}^l & l = 0, \dots, k-2 \\ (\frac{1}{n\alpha_{k-1}L} - \frac{1}{n\alpha_kL}) + \tau(1 - \frac{1}{n\alpha_{k-1}L}) & l = k-1 \\ \frac{1}{n\alpha_kL} & l = k \end{cases}$$

since

$$\begin{aligned} y_k &= x_k + \frac{1}{n\alpha_k L}(z_k - z_{k-1}) \\ &= \tau z_{k-1} + (1 - \tau)y_{k-1} + \frac{1}{n\alpha_k L}(z_k - z_{k-1}) \\ &= \tau z_{k-1} + (1 - \tau)(\sum_{l=0}^{k-2} c_{k-1}^l z_l + \frac{1}{n\alpha_{k-1}L}z_{k-1}) + \frac{1}{n\alpha_k L}(z_k - z_{k-1}) \\ &= (\sum_{l=0}^{k-2} (1 - \tau)c_{k-1}^l z_l) + ((\frac{1}{n\alpha_{k-1}L} - \frac{1}{n\alpha_kL}) + \tau(1 - \frac{1}{n\alpha_{k-1}L}))z_{k-1} + \frac{1}{n\alpha_k L}z_k \end{aligned}$$

As $\alpha_k \geq \alpha_{k-1}$, and $\alpha_0 = \frac{1}{nL}$, we have $c_k^l \geq 0$ for all l, k , and it is easy to check the coefficients sum to 1 for each k . \square

Lemma 4.9. Each iteration can be implemented in expected $O(N/n)$ time.

Proof. We show how to implement a iteration conditioned on i in time $O(\|A_{:,i}\|_0)$, where $\|A_{:,i}\|_0$ is the number of non-zeros in column i , thus give a expected running time of $O(N/n)$ for each iteration. We maintain the following quantities

$$z_k \in \mathbb{R}_{\geq 0}^n, az_k \in \mathbb{R}_{\geq 0}^m, y'_k \in \mathbb{R}^n, ay_k \in \mathbb{R}^m, B_{k,1}, B_{k,2} \in \mathbb{R}_+$$

with the following invariants always satisfied throughout the algorithm

$$Az_k = az_k \tag{8}$$

$$y_k = B_{k,1}z_k + B_{k,2}y'_k, \quad ay_k = B_{k,1}az_k + B_{k,2}ay_k \tag{9}$$

When $k = 0$, we let $az_k = Az_0, y'_k = y_0, ay_k = Ay_0, B_{k,1} = 0, B_{k,2} = 1$, and it is clear all the invariants are satisfied. For $k = 1, 2, \dots, T$:

- The step $x_k = \tau z_{k-1} + (1 - \tau)y_{k-1}$ does not need to be implemented.
- Computation of $\nabla_i f(x_k)$ requires the value of $p_j(x_k) = \exp(\frac{1}{\mu}(1 - (Ax_k)_j))$ for each j such that $A_{ji} \neq 0$, and we can get the value

$$(Ax_k)_j = \tau(Az_{k-1})_j + (1 - \tau)(Ay_{k-1})_j = (\tau + (1 - \tau)B_{k-1,1})(az_{k-1})_j + (1 - \tau)B_{k-1,2}ay_{k-1,j}$$

for each such j . This can be computed in $O(1)$ time for each j , and $O(\|A_{:,i}\|_0)$ time in total.

- The mirror descent step $z_k^{(i)} \stackrel{\text{def}}{=} \operatorname{argmin}_{z \in \Delta} \{V_{z_{k-1}}(z) + \langle z, n\alpha_k \xi_k^{(i)} \rangle\}$ is simply $z_k = z_{k-1} + \delta \mathbf{e}_i$ where $\delta \in \mathbb{R}$ can be computed in $O(1)$ time. $z_k = z_{k-1} + \delta \mathbf{e}_i$ yields $y_k = \tau z_{k-1} + (1 - \tau)y_{k-1} + \frac{\delta}{n\alpha_k L} \mathbf{e}_i$ by the gradient descent step. Therefore, we can update the values accordingly

$$z_k \leftarrow z_{k-1} + \delta \mathbf{e}_i, \quad az_k \leftarrow az_{k-1} + \delta A_{:,i}$$

and

$$\begin{aligned} B_{k,1} &\leftarrow \tau + (1 - \tau)B_{k-1,1} & B_{k,2} &\leftarrow (1 - \tau)B_{k-1,2} \\ y'_k &\leftarrow y'_{k-1} + \delta(-\frac{B_{k,1}}{B_{k,2}} + \frac{1}{n\alpha_k L} \frac{1}{B_{k,2}}) \mathbf{e}_i & ay_k &\leftarrow ay_{k-1} + \delta(-\frac{B_{k,1}}{B_{k,2}} + \frac{1}{n\alpha_k L} \frac{1}{B_{k,2}}) A_{:,i} \end{aligned}$$

We can verify that after the updates, the invariants still hold

$$\begin{aligned} y_k &= B_{k,1}z_k + B_{k,2}y'_k = B_{k,1}(z_{k-1} + \delta \mathbf{e}_i) + B_{k,2}(y'_{k-1} + \delta(-\frac{B_{k,1}}{B_{k,2}} + \frac{1}{n\alpha_k L} \frac{1}{B_{k,2}}) \mathbf{e}_i) \\ &= B_{k,1}z_{k-1} + B_{k,2}(y'_{k-1} + \delta(\frac{1}{n\alpha_k L} \frac{1}{B_{k,2}}) \mathbf{e}_i) \\ &= B_{k,1}z_{k-1} + B_{k,2}y'_{k-1} + \frac{\delta}{n\alpha_k L} \mathbf{e}_i \\ &= (\tau + (1 - \tau)B_{k-1,1})z_{k-1} + ((1 - \tau)B_{k-1,2})y'_{k-1} + \frac{\delta}{n\alpha_k L} \mathbf{e}_i \\ &= \tau z_{k-1} + (1 - \tau)y_{k-1} + \frac{\delta}{n\alpha_k L} \mathbf{e}_i \end{aligned}$$

It is also straightforward to verify $Ay_k = B_{k,1}az_k + B_{k,2}ay_k$ equals $Ay_k = \tau Az_{k-1} + (1 - \tau)Ay_{k-1} + \frac{\delta}{n\alpha_k L} A \mathbf{e}_i$. The updates are dominated by the updates on az_k and ay_k , which take $O(\|A_{:,i}\|_0)$ time.

□

References

- [1] Zeyuan Allen-Zhu and Lorenzo Orecchia. Nearly-linear time positive LP solver with faster convergence rate. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, STOC '15, pages 229–236, 2015. Newer version available at <http://arxiv.org/abs/1411.1124>.
- [2] Zeyuan Allen-Zhu and Lorenzo Orecchia. Using optimization to break the epsilon barrier: A faster and simpler width-independent algorithm for solving positive linear programs in parallel. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 1439–1456, 2015.
- [3] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(6):121–164, 2012.
- [4] Lisa Fleischer. A fast approximation scheme for fractional covering problems with variable upper bounds. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA 2004, New Orleans, Louisiana, USA, January 11-14, 2004, pages 1001–1010, 2004.
- [5] Christos Koufogiannakis and Neal E. Young. A nearly linear-time PTAS for explicit fractional packing and covering linear programs. *Algorithmica*, 70(4):648–674, 2014.

- [6] Michael Luby and Noam Nisan. A parallel approximation algorithm for positive linear programming. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 448–457, 1993.
- [7] Arkadi Nemirovski. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- [8] Yurii Nesterov. Smooth minimization of non-smooth functions. *Math. Program.*, 103(1):127–152, 2005.
- [9] Yurii Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- [10] Serge A. Plotkin, David B. Shmoys, and Éva Tardos. Fast approximation algorithms for fractional packing and covering problems. In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*, pages 495–504, 1991.
- [11] James Renegar. Efficient first-order methods for linear programming and semidefinite programming. *CoRR*, abs/1409.5832.
- [12] Neal E. Young. Sequential and parallel algorithms for mixed packing and covering. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 538–546, 2001.
- [13] Neal E. Young. Nearly linear-time approximation schemes for mixed packing/covering and facility-location linear programs. *CoRR*, abs/1407.3015, 2014.
- [14] Zeyuan Allen Zhu and Lorenzo Orecchia. Linear coupling: An ultimate unification of gradient and mirror descent. *CoRR*, abs/1407.1537, 2014.